Observations on using marks for pricing in multiclass packet networks to provide multidimensional QoS

Juan M. Alvarez and Bruce Hajek

University of Illinois at Urbana-Champaign

March 16, 2000

Outline

- Pricing and rate control.
- Bufferless model.
- Queueing model.
- Conclusions and further research.

Pricing and rate control

Goals for using pricing:

- Provide rate control in order to prevent congestion collapse in the network.
- Achieve an optimum use of network's resources under stable performance.

Assume: users have fixed expenditure rates.

In practice, such rates would be determined by user utility functions and performance obtained.

Consider the following user model:

$$x_s(n) = x_s(n-1) + \kappa\,(\; w_s - f_s(n-1)\;) \tag{1}$$

- $x_s(n)$: number of packets sent by user s in slot n.
- w_s: expenditure rate for user s.
- $f_s(n-1)$: charge imposed by the network to user s in slot $n-1$.

Note: $\frac{x_s(n)}{\kappa\ n} = w_s - \frac{f_s(1)+\ldots+f_s(n)}{n}$.

Bufferless model

- Single bufferless link.
- Discrete time with slot duration equal to one time unit.
- Each packet takes one slot to transmit.
- Link capacity: 10 packets per slot.
- Mark all packets that are transmitted in a congested slot.

Scenario 1:

- 40 users.
- Expenditure rate: 4.2 marks/slot.
- One packet class.
- Price per mark equal to 1.

Simulations show:

- Aggregate throughput: 8.152 packets/slot.
- Average loss: 11.95%.

Suppose half the users want low loss (say 1%) and are willing to sacrifice throughput.

Can the loss sensitive *users obtain their desired loss?*

Scenario 2:

- 20 throughput and 20 loss sensitive users.
- Expenditure rate for throughput sensitive users:
  2.1 marks/slot.
- Expenditure rate for loss sensitive users ranges from 0.021 to
  12 marks/slot.
- One packet class.
- Price per mark equal to 1.

Under these conditions it is not possible to satisfy the loss sensitive users QoS profile.

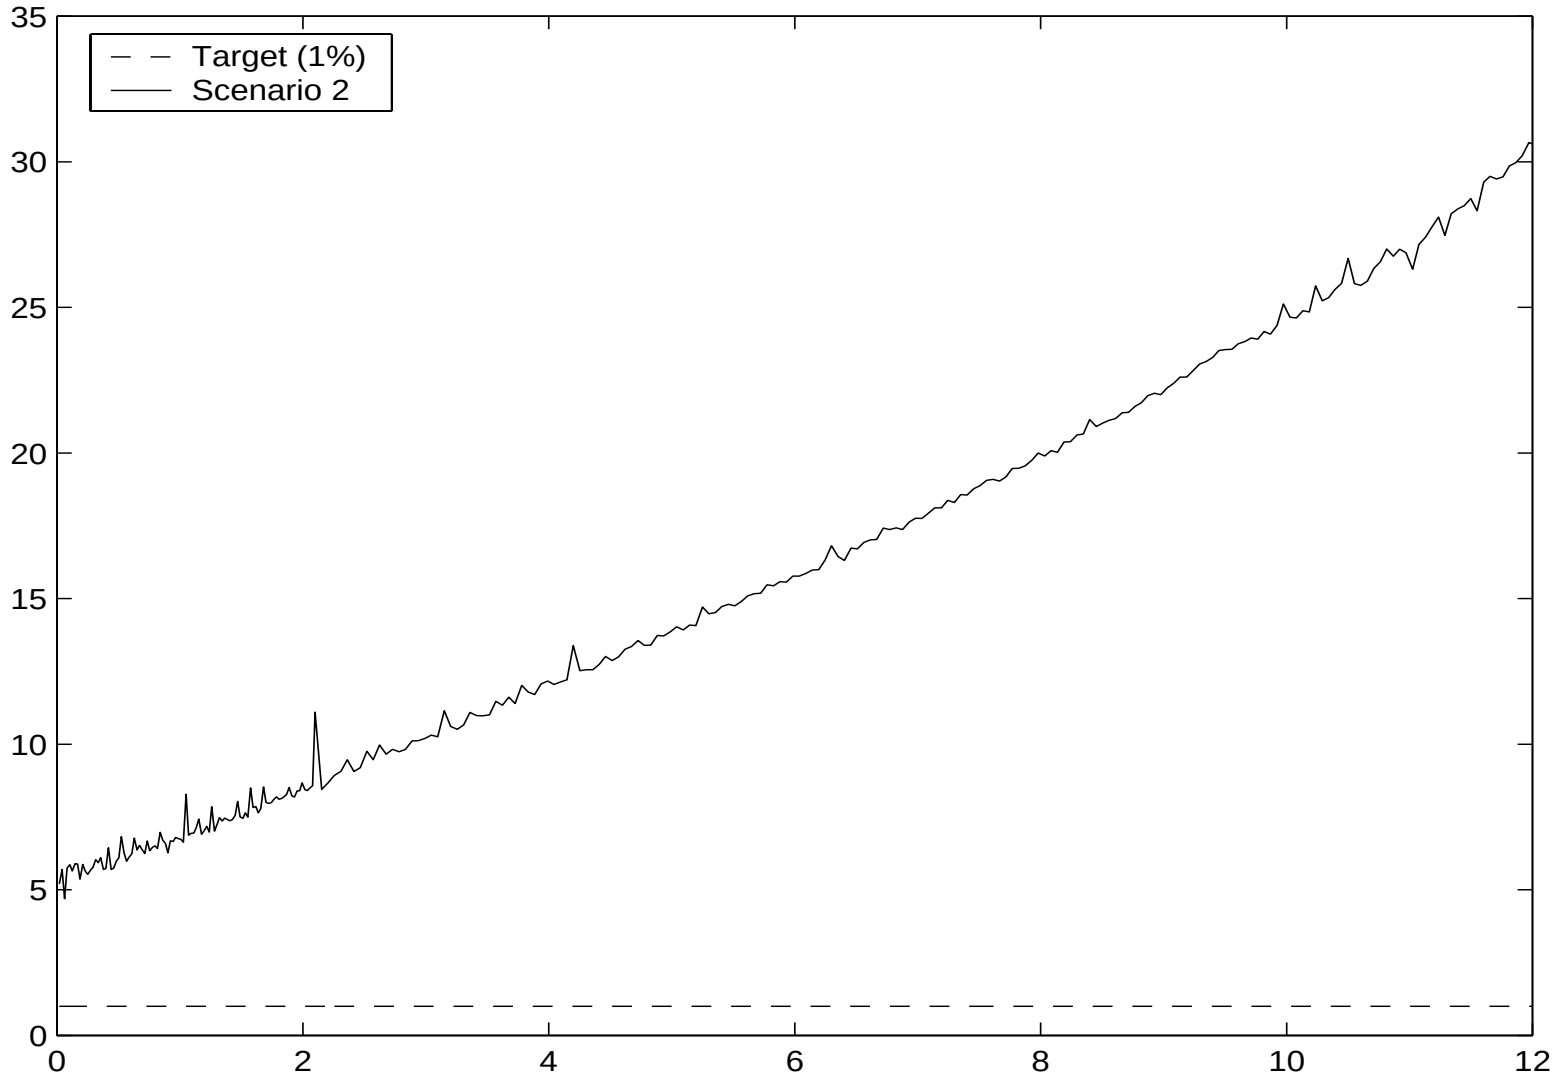
Target (1%)
Scenario 2
35
30
25
20
15
10
5
0
0
2
4
6
8
10
12

Solution: introduce multiple classes

- Allow two packet classes.
- Loss sensitive users declare all their packets as class 1.
- Throughput sensitive users declare all their packets as class 2.
- A portion of the capacity (c_1) is guaranteed for class 1 packets.

Sending rate update:

$$x_s(n) = x_s(n-1) + \kappa\,(\; w_s - p_s f_s(n-1)\;) \tag{2}$$

- $x_s(n)$: number of packets sent by user s in slot n.
- w_s: expenditure rate for user s.
- $f_s(n-1)$: number of feedback signals received by user s in slot $n-1$.
- p_s: price per mark for user s packet class.

Scenario 3:

- 20 throughput and 20 loss sensitive users.
- Expenditure rate for each type of user: 2.1 marks/slot.
- Two packet classes.
- Price per mark for class 1 packets ranges from 1 to 3.5.
- Price per mark for class 2 packets equal to 1.
- Resource provides service differentiation by using capacity reservation.

Under these conditions it is possible to satisfy the loss sensitive users QoS profile.

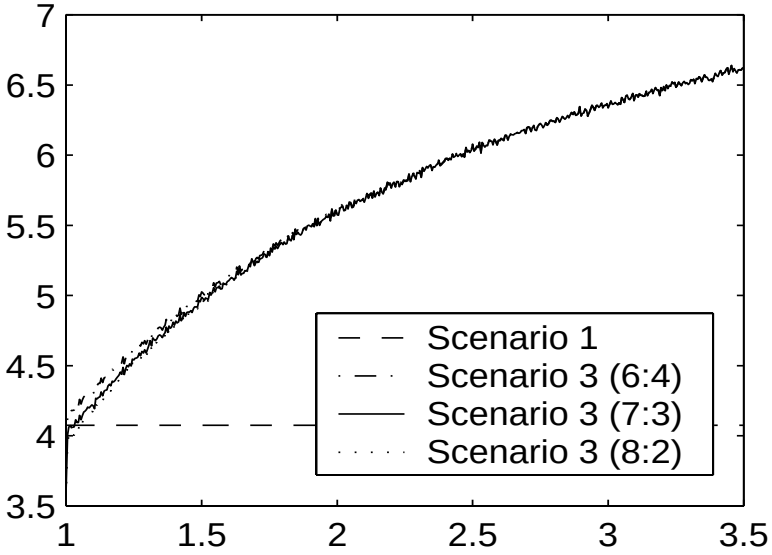
Scenario 1
Scenario 3 (6:4)
Scenario 3 (7:3)
Scenario 3 (8:2)

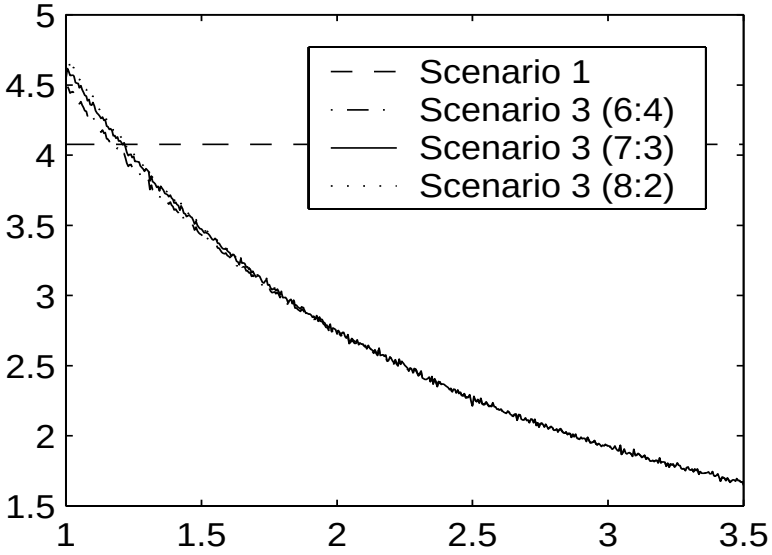
Scenario 1
Scenario 3 (6:4)
Scenario 3 (7:3)
Scenario 3 (8:2)

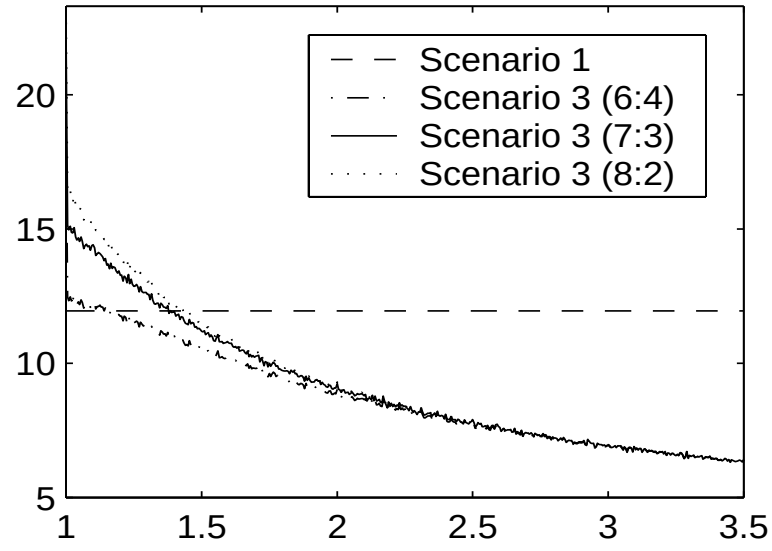
Scenario 1
Scenario 3 (6:4)
Scenario 3 (7:3)
Scenario 3 (8:2)

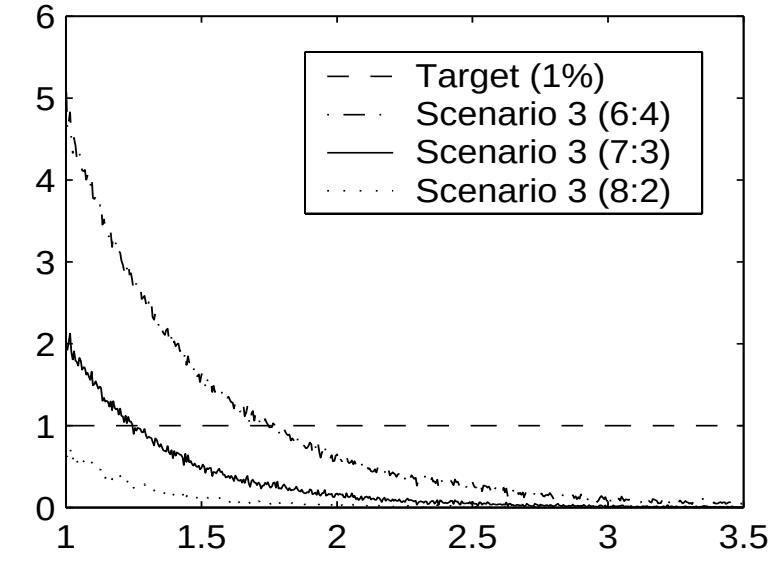
Target (1%)
Scenario 3 (6:4)
Scenario 3 (7:3)
Scenario 3 (8:2)

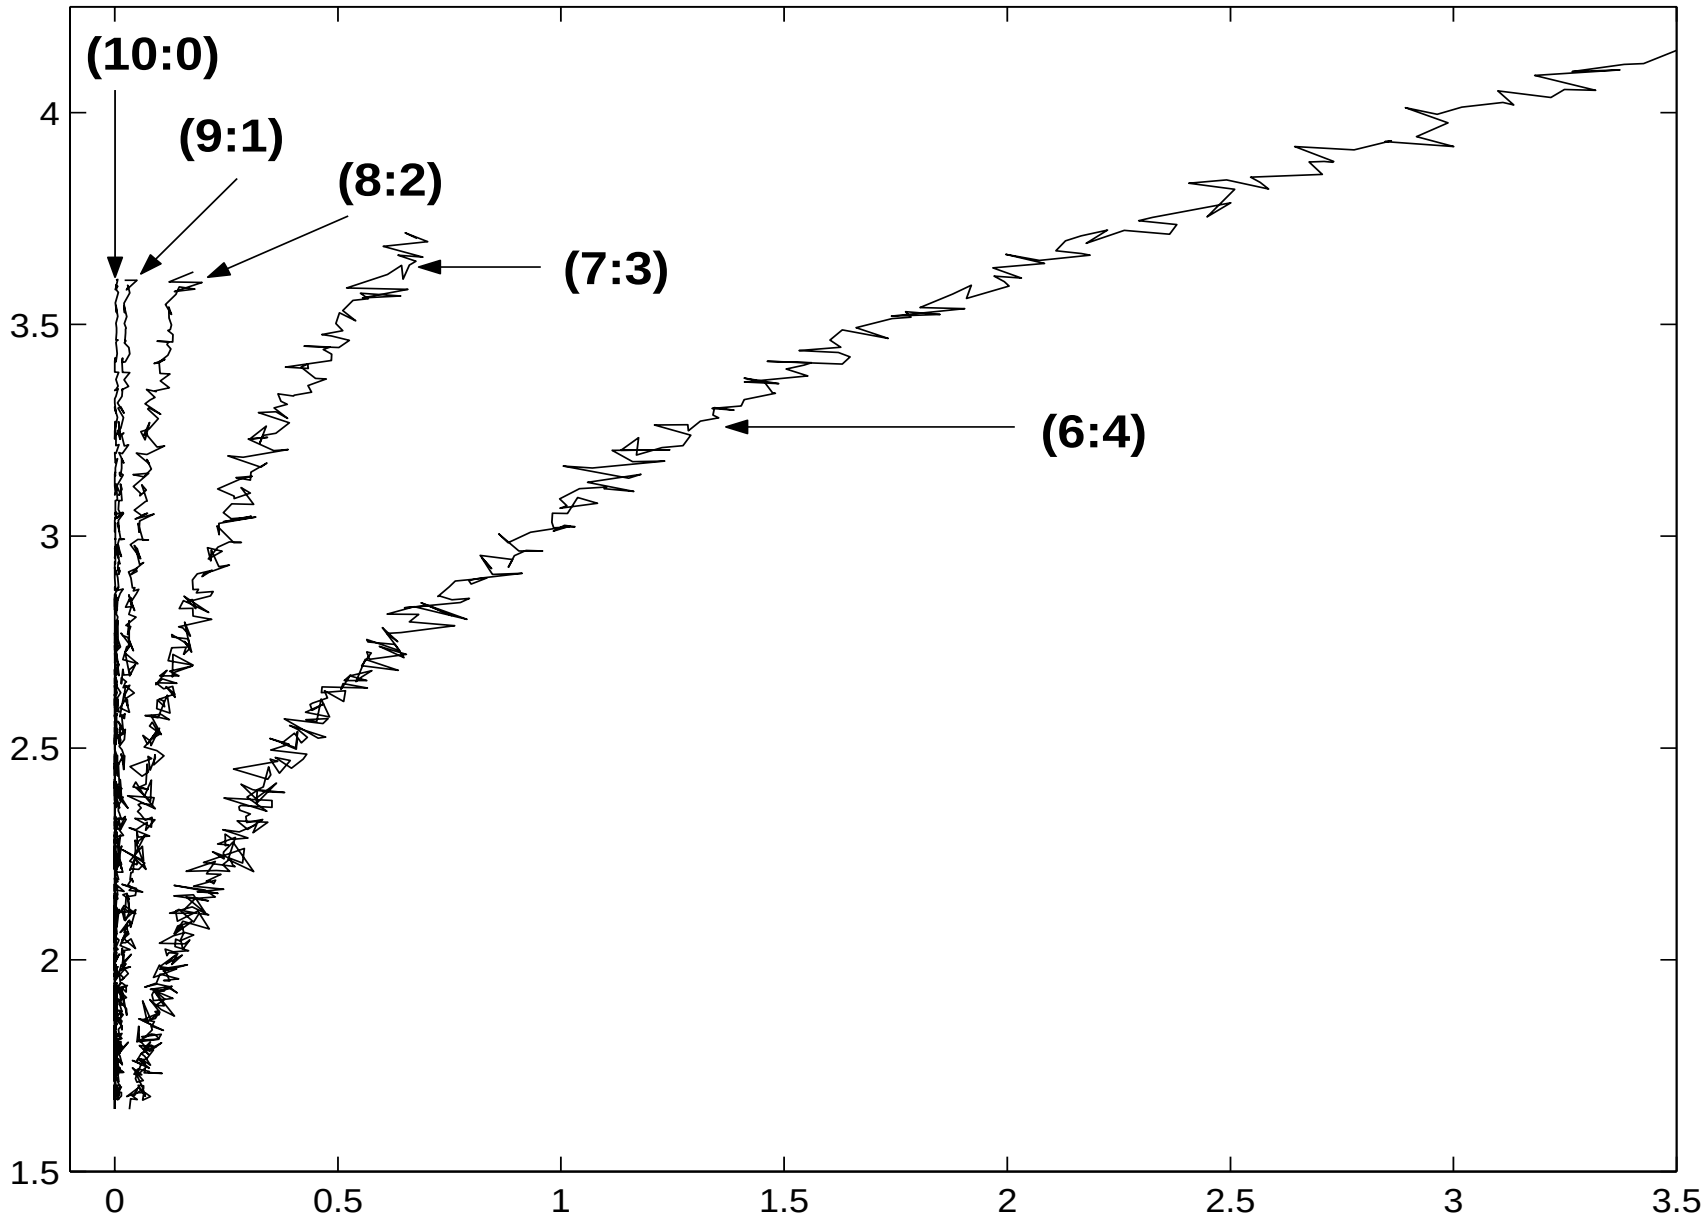
(10:0)
(9:1)
(8:2)
(7:3)
(6:4)
4
3.5
3
2.5
2
1.5
0
0.5
1
1.5
2
2.5
3
3.5

Queueing model

- Discrete time with slot duration equal to one time unit.
- Each packet takes one slot to transmit.
- Link can transmit one packet per slot.
- Buffer capacity: 10 packets.
- Packets are dropped at the time of their arrival if the buffer is full.
- Mark all packets after the first loss in a busy period, until the end of the same busy period.

Scenario 4:

- 40 throughput sensitive users.
- Expenditure rate: 0.42 marks/slot.
- Price per mark equal to 1.

Simulations show:

- Aggregate throughput: 0.92 packets/slot.
- Average loss: 3.95%
- Average delay: 5.87 slots.

Suppose again that half the users decide they want low loss (say 1%) and are willing to sacrifice throughput.

Can the loss sensitive *users obtain their desired loss?*

Scenario 5:

- 20 throughput and 20 loss sensitive users.
- Expenditure rate for throughput sensitive users:
  0.21 marks/slot.
- Expenditure rate for loss sensitive users ranges from 0.0021 to
  1.2 marks/slot.
- One packet class.
- Price per mark equal to 1.

Under these conditions it is not possible to satisfy the loss sensitive QoS profile.

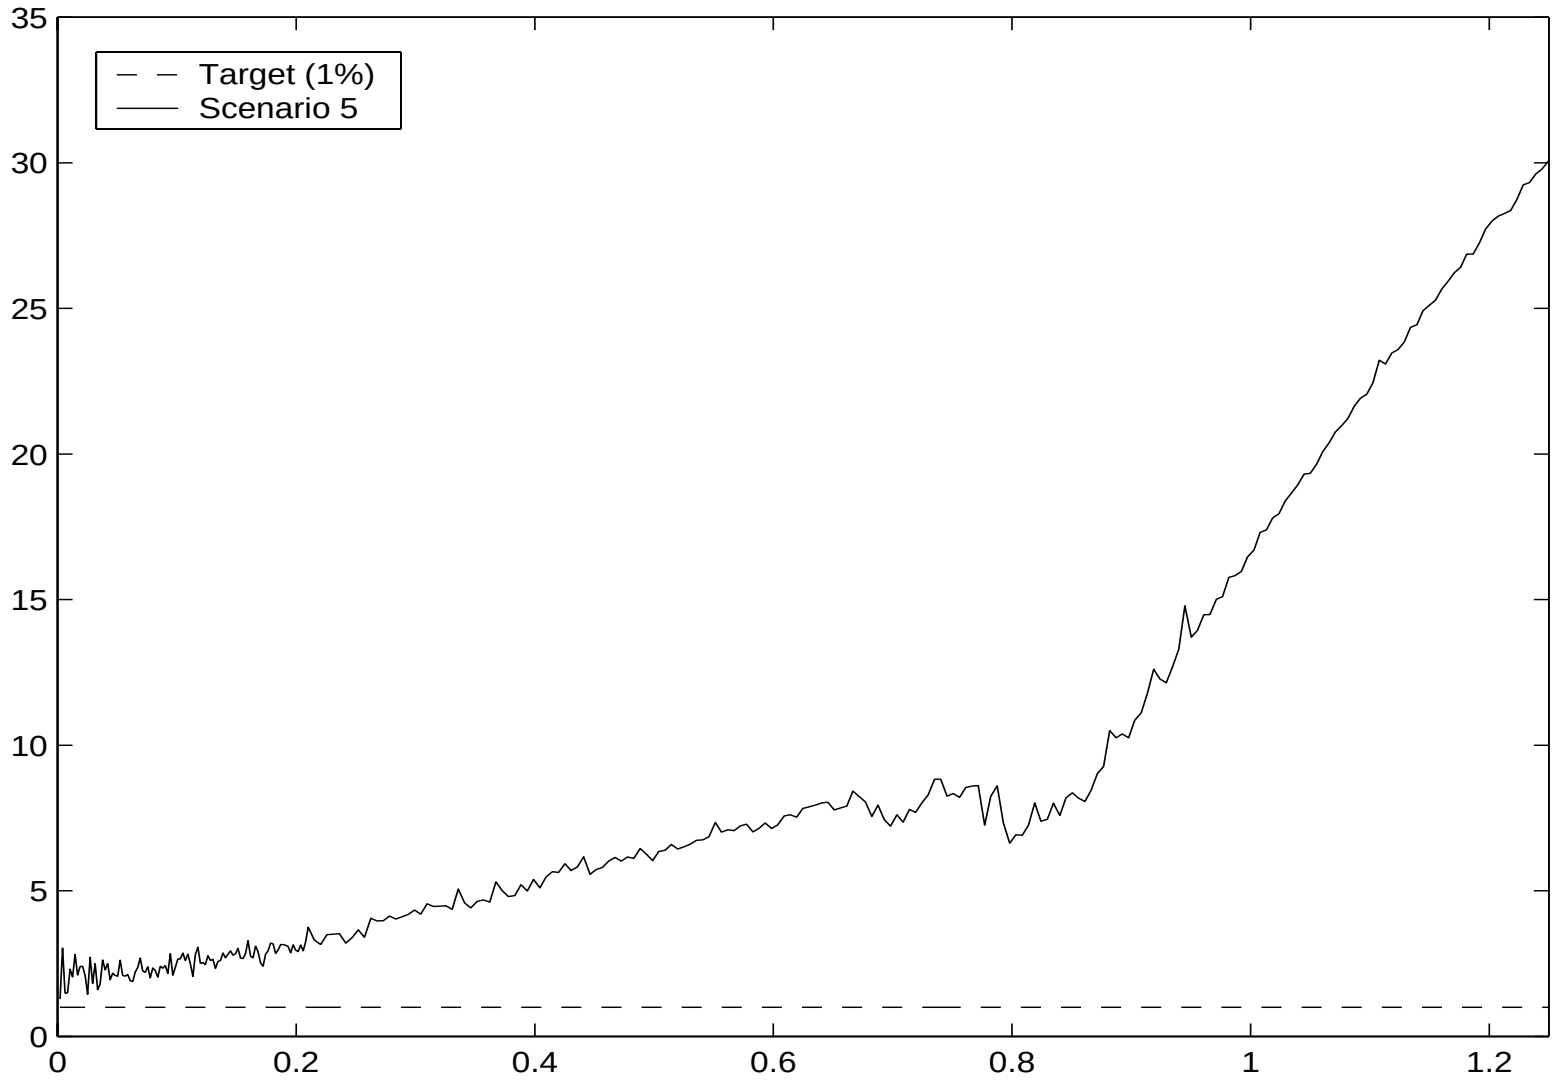
Target (1%)
Scenario 5
35
30
25
20
15
10
5
0
0
0.2
0.4
0.6
0.8
1
1.2

Solution: introduce multiple classes

- Allow two packet classes.
- Loss sensitive users declare all their packets as class 1.
- Throughput sensitive users declare all their packets as class 2.
- Resource provides differentiation in two ways:
  - A portion of the buffer (c_1) is guaranteed for class 1 packets.
  - Scheduling mechanisms:
    - Strict priority.
    - Weighted fair queueing (WFQ).

Scenario 6:

- 20 throughput and 20 loss sensitive users.
- Expenditure rate for each type of user: 0.21 marks/slot.
- Two packet classes.
- Price per mark for class 1 packets ranges from 1 to 7.
- Price per mark for class 2 packets equal to one unit.
- Resource provides service differentiation by using buffer reservation and scheduling.

Under these conditions it is possible to satisfy the loss sensitive QoS profile.

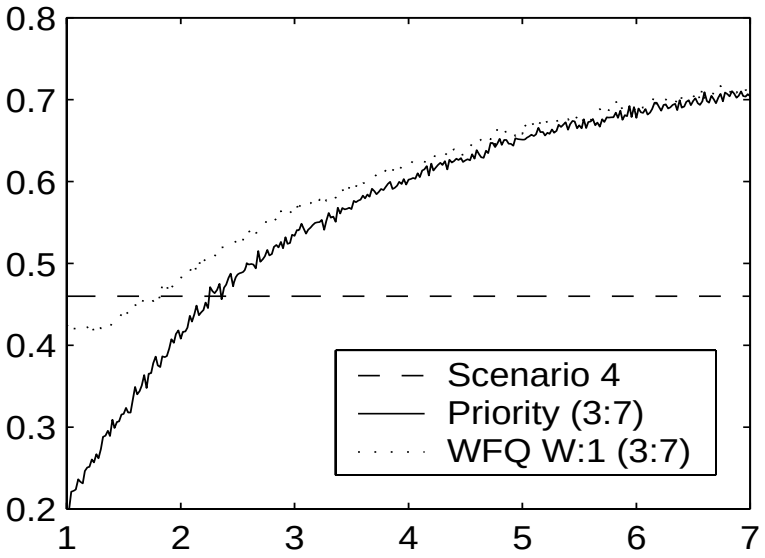
0.8
0.7
0.6
0.5
0.4
0.3
0.2
1
2
3
4
5
6
7
Scenario 4
Priority (3:7)
WFQ W:1 (3:7)

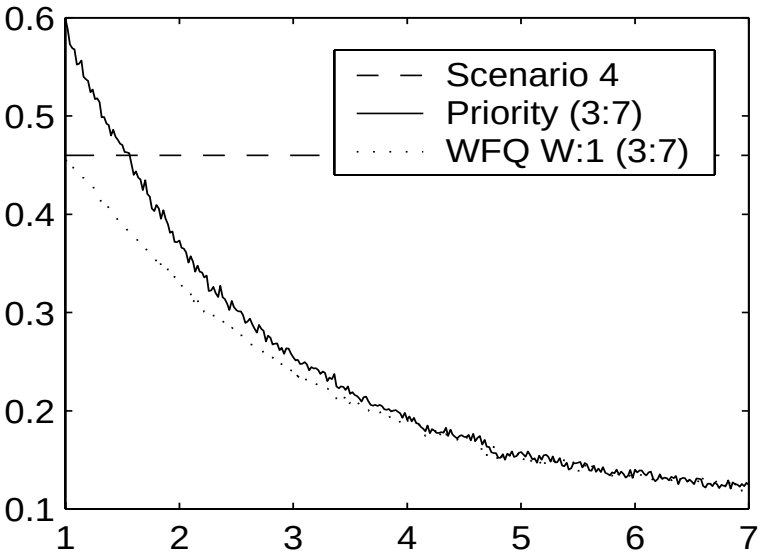
0.6
0.5
0.4
0.3
0.2
0.1
1
2
3
4
5
6
7
Scenario 4
Priority (3:7)
WFQ W:1 (3:7)

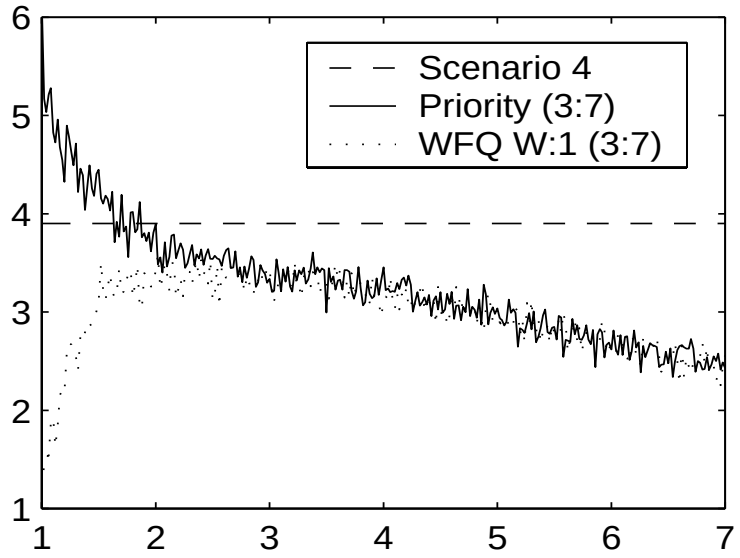
6
5
4
3
2
1
1
2
3
4
5
6
7
Scenario 4
Priority (3:7)
WFQ W:1 (3:7)

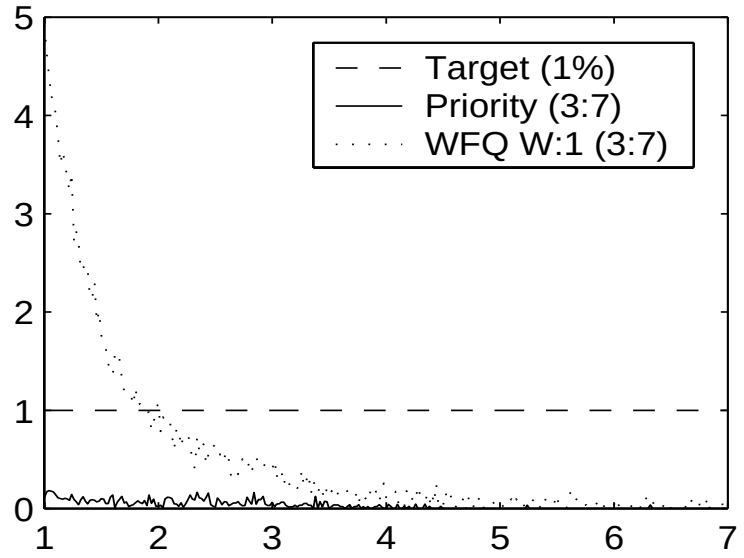
5
4
3
2
1
0
1
2
3
4
5
6
7
Target (1%)
Priority (3:7)
WFQ W:1 (3:7)

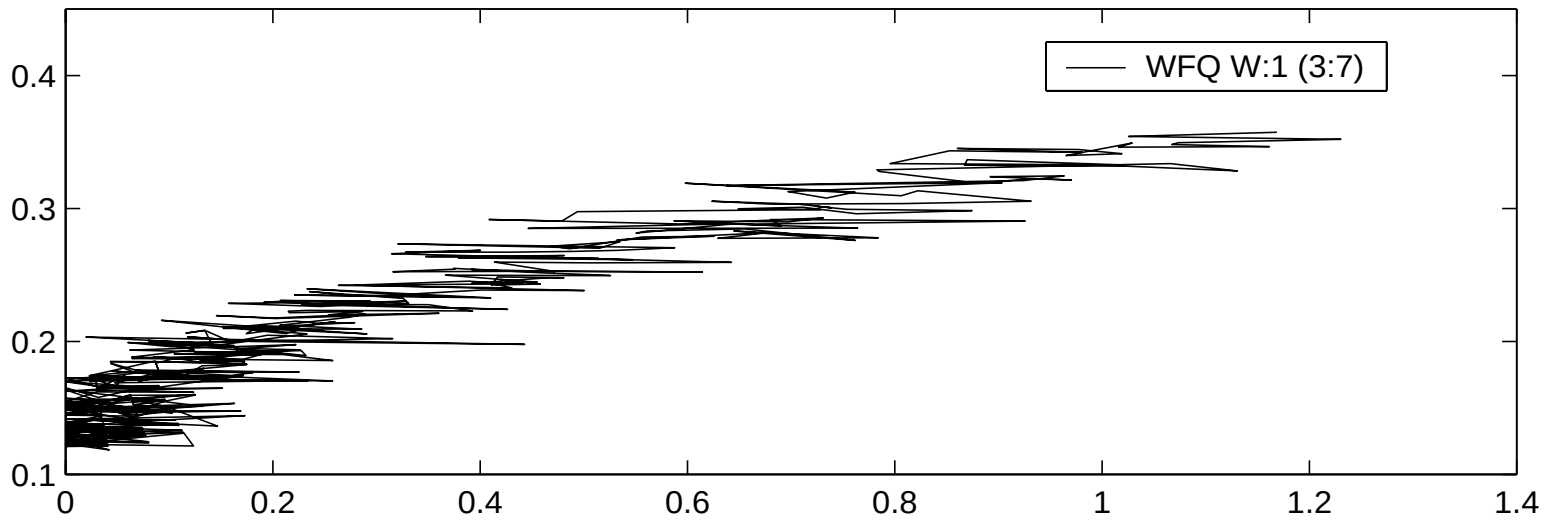
WFQ W:1 (3:7)
0.4
0.3
0.2
0.1
0
0.2
0.4
0.6
0.8
1
1.2
1.4

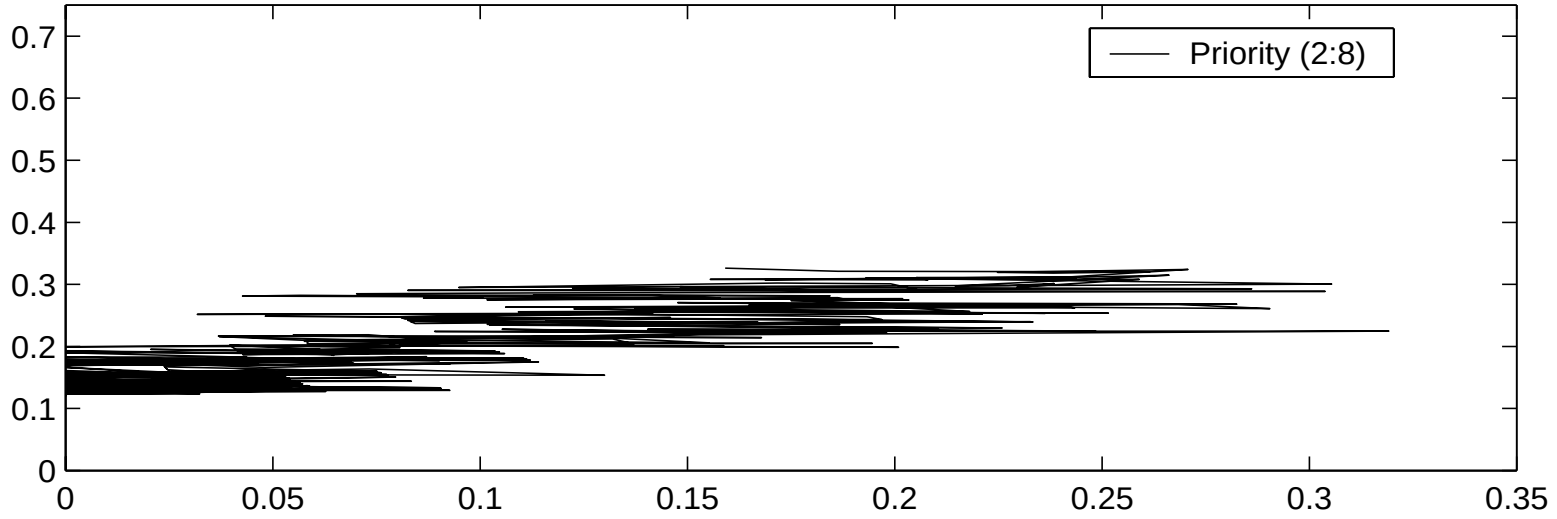
Priority (2:8)
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35

Conclusions

- The use of only one class of packets may not be sufficient to provide satisfactory multidimensional QoS performance.
- The use of multiple packet classes in conjunction with appropiate control mechanisms may be sufficient.
- The relaxation of some users' QoS parameters can be exploited by users that have tight requirements in those same parameters.

Further research

- Compatibility between different QoS requirements.
- Is QoS fundamentally multidimensional, or it can be reduced to one dimension?
- Compatibility of this rate control algorithm with existing ones (TCP,etc.).
- Stability if users can choose between classes frequently or vary their expenditure rate quickly.